Calibrating photometric redshifts of luminous red galaxies

Nikhil Padmanabhan,[1][*] Tamás Budavári,[2] David J. Schlegel,[3] Terry Bridges,[4] Jonathan Brinkmann,[5] Russell Cannon,[6] Andrew J. Connolly,[7] Scott M. Croom,[6] István Csabai,[8] Michael Drinkwater,[9] Daniel J. Eisenstein,[10] Paul C. Hewett,[11] Jon Loveday,[12] Robert C. Nichol,[13] Kevin A. Pimbblet,[9] Roberto De Propris,[14] Donald P. Schneider,[15] Ryan Scranton,[7] Uroš Seljak,[1] Tom Shanks,[16] István Szapudi,[17] Alexander S. Szalay[2] and David Wake[18]

[1]*Joseph Henry Laboratories, Jadwin Hall, Princeton University, Princeton, NJ 08544, USA*
[2]*Center for Astrophysical Sciences, Department of Physics & Astronomy, Johns Hopkins University, Baltimore, MD 21218, USA*
[3]*Department of Astrophysical Sciences, Peyton Hall, Princeton University, Princeton, NJ 08544, USA*
[4]*Physics Department, Queen's University, Kingston, ON, Canada, K7M 3N6*
[5]*Apache Point Observatory, 2001 Apache Point Road, Sunspot, NM 88349-0059, USA*
[6]*Anglo-Australian Observatory, PO Box 296, Epping, NSW 2121, Australia*
[7]*Department of Physics and Astronomy, University of Pittsburgh, 3941 O'Hara Street, Pittsburgh, PA 15260, USA*
[8]*Department of Physics, Eötvös University, Budapest, Pf. 32, Hungary, H-1518*
[9]*Department of Physics, University of Queensland, Qld 4072, Australia*
[10]*Steward Observatory, 933 N. Cherry Ave, Tucson, AZ 85721, USA*
[11]*Institute of Astronomy, Madingley Road, Cambridge CB3 0HA*
[12]*Astronomy Centre, University of Sussex, Falmer, Brighton BN1 9QJ*
[13]*Institute of Cosmology and Gravitation, University of Portsmouth, Portsmouth PO1 2EG*
[14]*Research School of Astronomy and Astrophysics, Australian National University, Weston, ACT 2611, Australia*
[15]*Department of Astronomy and Astrophysics, Pennsylvania State University, University Park, PA 16802, USA*
[16]*Department of Physics, University of Durham, South Road, Durham DH1 3LE*
[17]*Institute for Astronomy, 2680 Woodlawn Drive, Honolulu, HI 96822, USA*
[18]*Department of Physics, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA*

ABSTRACT

We discuss the construction of a photometric redshift catalogue of luminous red galaxies (LRGs) from the Sloan Digital Sky Survey (SDSS), emphasizing the principal steps necessary for constructing such a catalogue: (i) photometrically selecting the sample, (ii) measuring photometric redshifts and their error distributions, and (iii) estimating the true redshift distribution. We compare two photometric redshift algorithms for these data and find that they give comparable results. Calibrating against the SDSS and SDSS–2dF (Two Degree Field) spectroscopic surveys, we find that the photometric redshift accuracy is $\sigma \sim 0.03$ for redshifts less than 0.55 and worsens at higher redshift (~ 0.06 for $z < 0.7$). These errors are caused by photometric scatter, as well as systematic errors in the templates, filter curves and photometric zero-points. We also parametrize the photometric redshift error distribution with a sum of Gaussians and use this model to deconvolve the errors from the measured photometric redshift distribution to estimate the true redshift distribution. We pay special attention to the stability of this deconvolution, regularizing the method with a prior on the smoothness of the true redshift distribution. The methods that we develop are applicable to general photometric redshift surveys.

Key words: catalogues – surveys – galaxies: fundamental parameters.

1 INTRODUCTION

Since their inception, photometric redshifts (Koo 1985; Connolly et al. 1995; Gwyn & Hartwick 1996; Sawicki, Lin & Yee 1997; Hogg et al. 1998; Benítez 2000; Bolzonella, Miralles & Pelló 2000;

[*]E-mail: npadmana@princeton.edu

Csabai et al. 2000; Budavári et al. 2001; Collister & Lahav 2004) have provided a possible solution to the major limitation of large redshift surveys: i.e. that they are severely limited both in depth and area by the throughput of spectrographs. Photometric redshift algorithms essentially define a mapping from the observed photometric properties of galaxies to their redshifts and other physical properties such as luminosity and type. Given an accurate photometric redshift algorithm, one can map the observable Universe in three dimensions just by imaging in carefully chosen passbands. The relative efficiency of imaging compared with spectroscopy allows one to both go deeper and cover a larger area than is possible with traditional redshift surveys. Such surveys would be invaluable both for studies of large-scale structure, as well as understanding the evolution of galaxies. In addition, imaging surveys with well understood redshift distributions are essential for efforts to directly map the matter distribution using weak lensing.

Defining a photometric redshift catalogue involves fulfilling two requirements: one must photometrically specify a population of galaxies for which reliable photometric redshifts can be obtained and one must characterize the photometric redshift error distribution. Demonstrating this process is the purpose of this work, using the five-colour imaging of the Sloan Digital Sky Survey (SDSS; York et al. 2000) as an example.

Luminous red galaxies (LRGs) have long been recognized as a promising population for the application of photometric redshifts (Hamilton 1985; Gladders & Yee 2000; Eisenstein et al. 2001; Willis, Hewett & Warren 2001). These galaxies have remarkably uniform spectral energy distributions (SEDs; Schneider, Gunn & Hoessel 1983; Eisenstein et al. 2003) that are characterized by a strong break at 4000 Å caused by the accumulation of a number of metal lines. The redshifting of this feature through different filters gives these galaxies their characteristic red colours that are strongly correlated with redshift. This makes it easy to select these galaxies and to estimate photometric redshifts. In addition, these are among the most luminous galaxies in the Universe and map large cosmological volumes. Furthermore, LRGs are strongly correlated with clusters, making them an ideal tool for detecting and studying clusters. All of the above make LRGs an astrophysically interesting sample and an ideal candidate for a photometric redshift survey.

Measuring the photometric redshift error distribution requires a calibration set of spectroscopic redshifts that span a similar colour and magnitude range as the photometric catalogue. We use two redshift catalogues to calibrate the luminous red galaxy (LRG) photometric redshifts, the SDSS Data Release (DR) 1 (Abazajian et al. 2003)[1] LRG spectroscopic catalogue for redshifts <0.4 and the SDSS–2dF LRG spectroscopic catalogue (Cannon et al. 2003) for redshifts between 0.4 and 0.7. These catalogues have extremely good coverage of the LRG colour and magnitude selection criteria by design; the selection criteria we use have been strongly influenced by both these catalogues. In addition, we supplement the low-redshift catalogue with the SDSS Main galaxy catalogue complete to an *r*-band magnitude of 17.77.

A generic problem in interpreting analyses with photometric redshifts is estimating the conditional probability distribution, $P(z_{\mathrm{spectro}}|z_{\mathrm{photo}})$, as this allows us to connect the measurement (the photometric redshift) with the physical quantity (the actual redshift of the galaxy). This ability to connect photometric redshifts with actual

redshifts is essential to theoretically interpret results derived from photometric surveys and generically will be a significant source of systematic error. The simplest way to measure $P(z_{\mathrm{spectro}}|z_{\mathrm{photo}})$ is to directly measure it from a calibration data set. Unfortunately, $P(z_{\mathrm{spectro}}|z_{\mathrm{photo}})$ depends on the underlying redshift distribution and therefore, to obtain unbiased results, the calibration data and the actual data must sample the same redshift distribution. This is quite often not the case, because calibration data are drawn from heterogenous sources. We also note that simulations cannot solve this problem, because the $P(z_{\mathrm{spectro}}|z_{\mathrm{photo}})$ derived will depend on the simulated redshift distribution, which might differ significantly from the true distribution.

The approach that we favour in this paper is to use the Bayes theorem to relate $P(z_{\mathrm{spectro}}|z_{\mathrm{photo}})$ to $P(z_{\mathrm{photo}}|z_{\mathrm{spectro}})$, using the true redshift distribution, dN/dz of the photometric sample. For samples selected only with an apparent magnitude cut, one can estimate dN/dz directly from the galaxy luminosity function (for e.g. Budavári et al. 2003). This approach is significantly harder for samples, like the ones considered in this paper, that involve multiple magnitude and colour cuts, as it involves the joint luminosity–colour distribution functions that are poorly understood.

We present an alternative method to estimate dN/dz in this paper, that starts from the observation that the observed photometric redshift distribution is just the true redshift distribution convolved with the photometric redshift errors. Phrased as such, estimating dN/dz is simply the problem of deconvolving the redshift errors from the measured redshift distribution. This problem, like all deconvolution problems, is ill-conditioned and must be regularized to obtain a stable solution.

This paper is organized as follows. Section 2 describes the two sources for our calibration data, the SDSS and SDSS–2dF surveys, and presents our selection criteria for LRGs. In Section 3, we describe two photometric redshift algorithms and calibrate them against the catalogues from the previous section and measure the photometric redshift error distribution. Section 4 discusses using this error distribution to invert the observed photometric redshift distribution to reconstruct the true redshift distribution, while Section 5 summarizes our conclusions. Whenever necessary, we have assumed a cosmology with $\Omega_{\mathrm{m}} = 0.3$, $\Omega_{\Lambda} = 0.7$ and $H_0 = 100\,h\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$.

2 SELECTING RED GALAXIES

We start by describing the data that form our calibration data set, the spectroscopic (Main and Luminous Red Galaxy, LRG) survey of the SDSS and the SDSS–2dF LRG survey; the reader is referred to the appropriate technical documents (Eisenstein et al. 2001; Strauss et al. 2002) for a more detailed description. We then present the exact cuts used to construct our sample of LRGs. These are similar in spirit to those in Eisenstein et al. (2001) although they differ in detail.

Because the photometry for the two catalogues we use is from the SDSS, we restrict our discussion in this paper to the SDSS five-filter photometric system (Fukugita et al. 1996; Smith et al. 2002). The methods can be generalized to an arbitrary photometric system. Except where explicitly specified, we will use SDSS model magnitudes (Stoughton et al. 2002); for instance, *g* will refer to an SDSS *g*-band model magnitude. SDSS Petrosian magnitudes will be denoted by the subscript Petro, e.g. r_{Petro} is the SDSS *r*-band Petrosian magnitude.

Finally, a comment on the magnitude system used: it has become traditional to use AB magnitudes (Oke & Gunn 1983) for

---

[1] We note that DR1 here only refers to the area coverage; the reduction pipelines used are identical to those for DR2 and DR3. In particular, the model magnitude bug in the DR1 reductions does not affect this paper.

estimating photometric redshifts. The SDSS magnitudes are close to AB magnitudes, but differ at the millimagnitude level (Abazajian et al. 2004). The final zero-point corrections for the SDSS have yet to be determined; we use the best estimate of these offsets available at the time of writing. The offsets applied are $\Delta(u, g, r, i, z) = (-0.042, 0.036, 0.015, 0.013, -0.002)$. We note that the photometric redshifts are not very sensitive to the precise values of these offsets; not including them changes the measured redshifts by $\Delta z \sim 0.005$, completely subdominant to the photometric redshift errors.

2.1 The SDSS surveys

The SDSS is an ongoing survey to image approximately π sr of the sky and follow up approximately one million of the detected objects spectroscopically. The imaging is carried out by drift-scanning the sky (Gunn et al. 1998) in photometric conditions (Hogg et al. 2001), in five (*ugriz*) bands using a specially designed wide-field camera. Using these imaging data as a source, objects targeted for spectroscopy (Strauss et al. 2002; Blanton et al. 2003) are observed with a 640-fibre spectrograph on the same telescope. All of these data are processed by completely automated pipelines that detect and measure photometric properties of objects, and astrometrically calibrate the data (Pier et al. 2003; Lupton, in preparation). The SDSS is close to completion and has had three major data releases (Early Data Release, EDR, Stoughton et al. 2002; DR1, Abazajian et al. 2003; DR2, Abazajian et al. 2004). This paper will limit itself to DR1, with approximately 168 000 spectra.

The data used in this paper are from the Main (Strauss et al. 2002) and LRG (Eisenstein et al. 2001) surveys. The Main galaxy sample is a magnitude limited survey targeting all galaxies with $r_{\text{Petro}} < 17.77$. The SDSS LRG sample targets a smaller set of galaxies with $r_{\text{Petro}} < 19.5$; these galaxies are colour selected to have strong 4000-Å breaks allowing a spectroscopic determination of their redshifts even though they are ~ 2 mag fainter than the Main galaxy sample. The selection methodology of these galaxies forms the basis both of the SDSS–2dF survey, which we now discuss, and the selection criteria we present in Section 2.3.

2.2 The SDSS–2dF survey

The second set of observations are the first data obtained as part of the SDSS–2dF LRG survey. This redshift survey, started in early 2003, exploits the marriage of two facilities: the wide-angle, multicolour, imaging data of the SDSS and the 2dF spectrograph on the 4-m Anglo-Australian Telescope (AAT; Lewis et al. 2002). The SDSS–2dF LRG survey is being carried out in tandem with the SDSS–2dF quasi-stellar object (QSO) survey to ensure optimal use of the 400 spectroscopic fibres available in the 2dF spectrograph.

The goal of the SDSS–2dF LRG survey is to replicate the selection of SDSS LRGs but at a higher redshift, by going to fainter apparent luminosities. In particular, we aim to closely match the space density, luminosity range and colours of the lower-redshift SDSS LRGs, thus allowing the study of the evolution of a single population of massive galaxies over a large redshift range. To achieve this goal, we use the same methodology as outlined in Eisenstein et al. (2001) for selecting the low-redshift SDSS LRGs, but adapt the colour and magnitude cuts to preferentially select LRGs in the redshift range $0.45 < z < 0.7$. The SDSS–2dF LRG cuts we use are similar to those of the Cut II SDSS LRG sample discussed in detail in Eisenstein et al. (2001). However, because of the larger telescope (AAT) and the longer integration times possible, we relax the $r < 19.5$ magnitude limit of Cut II (which resulted in a severe redshift limit of $z \simeq 0.45$

for the SDSS LRGs) to $i \leqslant 20$. As discussed in Eisenstein et al. (2001), the selection of LRGs above $z \simeq 0.4$ is actually easier than selecting them at lower redshifts because the 4000-Å break moves into the SDSS r band and, therefore, the SDSS $r - i$ colour is an effective estimate of the redshift, while the $g - r$ colour is a proxy for the rest-frame colour of the galaxy.

The details of the SDSS–2dF LRG selection criteria will be presented elsewhere. However, as shown in Nichol (2003) and Fig. 3, the SDSS–2dF selection criteria successfully reproduce the luminosity range covered by the lower-redshift SDSS LRGs (Eisenstein et al. 2001, both Cut I and Cut II,) over the expected range of redshifts from $0.4 < z < 0.75$. Note that the $r - i$ colour selection is very effective at isolating high-redshift galaxies, with 90 per cent of the galaxies having redshifts between 0.4 and 0.7, and virtually none with redshifts <0.3. The SDSS–2dF LRG and QSO surveys are underway with the goal of obtaining the final sample of $\simeq 10\,000$ LRGs and quasars. The data we use are all the data observed through 2003 with reliable spectroscopic redshifts, a sample of ~ 3000 galaxies.

2.3 Selection criteria

We now discuss the construction of a photometric sample of LRGs. Although the selection criteria we present here (including the terminology) are based on the spectroscopic selection used to construct the two samples discussed above, we emphasize that these are not the specific selection criteria for either sample, but rather are a synthesis of different selection techniques. The goal of these selection criteria is to photometrically select a uniform sample of LRGs over the redshift range $0.2 < z < 0.7$.

Fig. 1 shows a model spectrum of an early-type galaxy from the stellar population synthesis models of Bruzual & Charlot (2003). This particular spectrum is derived from a single burst of star formation 11 Gyr ago (implying a redshift of formation, $z_{\text{form}} \sim 2.6$), evolved to the present, and is typical of LRG spectra. In particular, the 4000-Å break is very prominent. In order to motivate our selection criteria, we passively evolve this spectrum in redshift (in
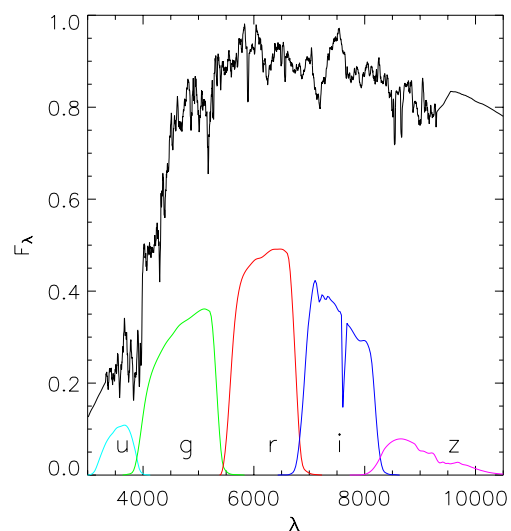

Figure 1. A model spectrum of an early-type galaxy from (Bruzual & Charlot 2003). The model was formed from a single burst of star formation 11 Gyr ago and assumes a solar metallicity. Note the prominent break in the spectrum at 4000 Å. Also overplotted are the response functions (including atmospheric absorption) for the Sloan Digital Sky Survey (SDSS) filters.
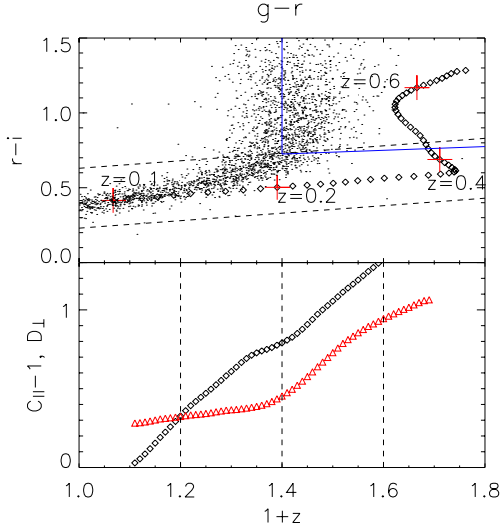
Figure 2. The top panel shows simulated $g - r$ and $r - i$ colours of an early-type galaxy as a function of redshift. The spectrum used to generate the track is the same as in Fig. 1, but evolved in redshift. Also shown are the colour cuts for Cut I (dashed, black) and Cut II galaxies (solid, blue in electronic figure). The points show the stellar locus as determined by a sample of stars with r-band magnitudes less than 19.5. The lower panel shows the colours $c_{||}$ (diamonds, black) and $d_\perp$ (triangles, red in electronic figure), as a function of redshift. Also shown are fiducial redshift boundaries for Cut I (0.2–0.4) and Cut II (0.4–0.6). Note that the range in $g - r$ is identical to the range in $1 + z$.

particular, taking the evolution of the strength of the 4000-Å break into account) and project it through the SDSS filters; the resulting colour track in g–r–i space as a function of redshift is shown in Fig. 2. The bend in the track around $z \sim 0.4$, caused by the red-shifting of the 4000-Å break from the g to r band, naturally suggests two selection criteria: a low-redshift sample (Cut I), nominally from $z \sim 0.2$–0.4, and a high-redshift sample (Cut II), from $z \sim 0.4$–0.6. We define the two colours (Eisenstein et al. 2001; Eisenstein, private communication)

$$c_\perp \equiv (r - i) - (g - r)/4 - 0.18, \tag{1}$$

$$d_\perp \equiv (r - i) - (g - r)/8 \approx r - i. \tag{2}$$

We now make the following colour selections:

$$\text{Cut I} = |\, c_\perp \,| < 0.2; \tag{3}$$

$$\text{Cut II} = d_\perp > 0.55, \tag{4}$$

$$g - r > 1.4; \tag{5}$$

as shown in Fig. 2. The final cut, $g - r > 1.4$, isolates our sample from the stellar locus. In addition to these selection criteria, we eliminate all galaxies with $g - r > 3$ and $r - i > 1.5$; these constraints eliminate no real galaxies, but are effective in removing stars with unusual colours.

Unfortunately, as emphasized in Eisenstein et al. (2001), these simple colour cuts are not sufficient to select LRGs as a result of an accidental degeneracy in the SDSS filters that causes all galaxies, irrespective of type, to lie very close to the low-redshift early-type locus. We therefore follow the discussion there and impose a cut in absolute magnitude. We implement this by defining a colour to use as a proxy for redshift and then translating the absolute magnitude cut

into a colour–magnitude cut. We see from Fig. 2 that $d_\perp$ correlates strongly with redshift and is appropriate to use for Cut II. For Cut I, we define

$$c_{||} = 0.7(g - r) + 1.2(r - i - 0.18), \tag{6}$$

which is approximately parallel to the low-redshift locus. Given these, we further impose:

$$\text{Cut I} = r_{\text{Petro}} < 13.6 + c_{||}/0.3,$$
$$r_{\text{Petro}} < 19.7; \tag{7}$$

$$\text{Cut II} = i < 18.3 + 2d_\perp,$$
$$i < 20. \tag{8}$$

Note we use r_{Petro} for consistency with the SDSS LRG target selection. We note that Cut I is identical (except in the numerical values of the magnitude cuts in equation 7) to the SDSS LRG Cut I, while the numerical values for Cut II were chosen to yield a population consistent with Cut I (see below). This was intentionally done to maximize the overlap between any sample selected using these cuts and the SDSS LRG spectroscopic sample. The switch to the i band for Cut II also requires some explanation. As is clear from Fig. 1, the 4000-Å break is moving through the r band throughout the fiducial redshift range of Cut II. This implies that the K-corrections to the r band are very sensitive to redshift; the i-band K-corrections are much less sensitive to redshift allowing for a more robust selection.

The results of applying these cuts to the spectroscopic catalogues are shown in Fig. 3. Because the SDSS spectroscopic catalogue is at low redshift, we trim the catalogue using Cut I, while the higher-redshift SDSS–2dF data are trimmed with Cut II. Our calibration data set has 45 744 Cut I galaxies and 1474 Cut II galaxies. The large number of low-redshift galaxies that pass Cut I indicate a failure of our selection criteria at redshifts lower than $z \sim 0.15$, as already noted by Eisenstein et al. (2001). We however leave these galaxies
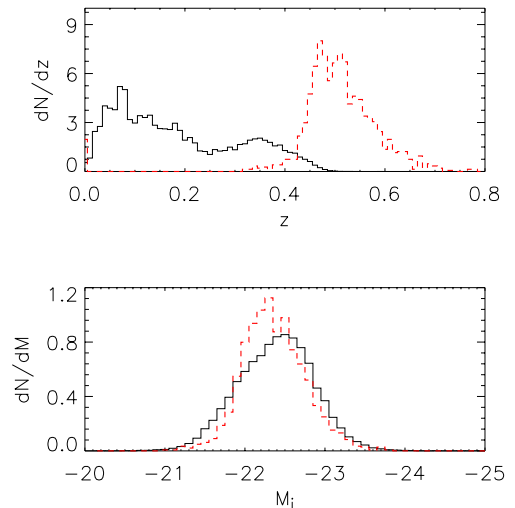

Figure 3. The top panel shows the spectroscopic redshift distribution, $\mathrm{d}N/\mathrm{d}z$, of the Sloan Digital Sky Survey (SDSS; solid, black) and the SDSS–2dF (dashed, red in electronic figure) samples trimmed using the selection criteria of Section 2.3. Note that the SDSS sample is dominated by the low-redshift main sample, accounting for the low normalization at high redshift. The lower panel shows the i-band absolute magnitude distribution for the two samples, demonstrating that our absolute magnitude cuts are selecting a sample with $M_i \sim -22$ as desired. Both $\mathrm{d}N/\mathrm{d}z$ and $\mathrm{d}N/\mathrm{d}M$ are normalized so that they integrate to unity.

in our analysis, because they will contaminate any photometrically selected sample and it will be necessary to understand their photometric redshift distributions. No such problem exists for the Cut II galaxies, which have a negligible fraction of $z < 0.4$ galaxies. The most significant contaminant for Cut II are M stars. The $g - r > 1.4$ cut removes most of these, although there is a small residual level of contamination. Analyses using this or similar samples will have to estimate the effect of this contamination on their results.

The lower panel of Fig. 3 shows the absolute magnitude distribution of both Cut I and Cut II galaxies. As expected, the colour magnitude cuts restrict the sample to bright galaxies; the median i-band magnitude is $M_i \sim -22$. Note that the Cut I and Cut II galaxies probe similar luminosities. The Cut I magnitude distribution also has a tail extending to low luminosities; this is the failure of the selection criteria at low redshift we encountered above.

3 PHOTOMETRIC REDSHIFT ESTIMATION

Photometric redshift estimation techniques can be classified into two groups, empirical and template-fitting methods. Empirical methods (Connolly et al. 1995; Wang, Bahcall & Turner 1998; Brunner, Connolly & Szalay 1999) are based on the observational fact that galaxies are restricted to a low dimensional surface in the space of their colours and redshift. Using a training set of galaxies, these methods attempt to parametrize this surface, either with low-order polynomials, nearest-neighbour searches or neural networks. The advantage of these methods is that they attempt to measure these relationships directly from the data and so implicitly correct for any calibration biases present. A publicly available example is the Artificial Neural Network code (ANNZ; Firth, Lahav & Somerville 2003; Collister & Lahav 2004) that trains a neural network to learn the relation between photometry and redshift from an appropriate training set of observed galaxies whose redshifts are already known. This code has a photometric redshift accuracy similar to the methods described below(Collister & Lahav 2004).

The fact that these methods rely on training sets is their greatest disadvantage. For these methods to work, the training set must densely sample the entire redshift–colour space of interest, as it is difficult to extrapolate outside the domain of the training set. Most training sets, including the samples constructed above, violate the above requirement and therefore are of limited utility. Template-based methods do not suffer from these drawbacks and form the basis for the two algorithms used in this paper, which we now discuss.

3.1 Simple template fitting

Template-fitting methods start with a set of model spectra (the templates) of galaxies, either from spectrophotometrically calibrated observations of galaxies (Coleman, Wu & Weedman 1980) or from stellar synthesis models (Le Borgne & Rocca-Volmerange 2002; Bruzual & Charlot 2003). These methods then attempt to reconstruct the observed colours of galaxies by some (appropriately redshifted) linear combination of the templates, projected through the appropriate filters. The best-fitted redshift is then an estimate of the true redshift of the galaxy. Concretely, if we denote the templates by $\Psi^i(z)$, this algorithm can be cast as a minimization of

$$\chi^2(c_i, z) = \sum_\alpha \left\{ \frac{f_\alpha - R_\alpha \left[\sum_i c_i \Psi^i(z) \right]}{\sigma_\alpha} \right\}^2 , \qquad (9)$$

where f_α is the observed flux (with error σ_α) of the galaxy in the α filter and $R_\alpha(\Psi)$ projects the spectrum Ψ on to the α filter. For definiteness, we work with the AB photometric system (Oke & Gunn 1983), where the apparent magnitude of a galaxy, m_{AB}, is related to its spectral energy distribution (SED), Ψ (with units of W m^{-2} Hz^{-1}), by

$$m_{AB} = -2.5 \log_{10} \left[\frac{\int d\nu \, \nu^{-1} \Psi(\nu) W_\alpha(\nu)}{\int d\nu \, \nu^{-1} g(\nu) W_\alpha(\nu)} \right] , \qquad (10)$$

where W_α is the response of the α filter. The reference SED is given by $g(\nu) = 3631$ Jy (where 1 Jy $= 10^{-26}$ W m^{-2} Hz^{-1}).

One of the advantages of the LRGs is that their spectra are well described by a single template (Eisenstein et al. 2003). We find that the LRG colours are well described by a Bruzual & Charlot (2003)[2] single instantaneous burst template at solar metallicity with the burst occurring when the Universe was 2.5 Gyr old. The template evolves with time, becoming redder and increasing the 4000-Å break, as the more massive hot stars die. To incorporate this evolution, we interpolate between models with bursts with ages of [11, 5, 2.5, 1.4, 0.9, 0.64, 0.1] Gyr to calculate the template as a function of redshift. The photometric redshifts we derive are insensitive to the precise time of the burst and, therefore, we do not attempt any optimization of this parameter.

The implementation of this method we use is part of the IDL-SPEC2D SDSS spectroscopic reduction pipeline (Burles & Schlegel, in preparation) and can be downloaded through the World Wide Web (WWW).[3]

3.2 Hybrid methods

Obviously, this simple template-fitting algorithm is effective only when the templates accurately describe the photometric properties of the galaxies for which one wants to estimate redshifts. This suggests generalizing the template-fitting algorithm by incorporating features of empirical methods (Budavári et al. 2000; Csabai et al. 2000; Budavári et al. 2001). The basic approach is to divide a training set into spectral classes corresponding to each of the templates. Given these training sets for the individual templates, one can repair the templates by adjusting them to better reproduce the measured colours of the galaxies in the training set. By repeating this classification and repair procedure, one can obtain an improved template set that yields more reliable photometric redshifts (Csabai et al. 2003). Moreover, this process of adjusting the templates to agree with observations makes hybrid methods potentially less sensitive to potential systematic problems resulting from errors in the filter curves or photometric zero-points. We refer the reader to the above papers for details on the implementation of this algorithm. For the LRGs, we start with an elliptical template from Coleman et al. (1980) and apply the above algorithm to optimize it. This is done in an iterative manner and converges after typically three iterations.

This single optimized template is used for an initial redshift estimate for all LRGs. The SDSS LRG sample is, however, selected assuming a passively evolving elliptical template (Eisenstein et al. 2001). Therefore, we expect to gain in photometric redshift accuracy if we allow the LRG spectral template to evolve with redshift. The SDSS and SDSS–2dF redshift samples are subdivided into three redshift intervals $0.00 < z < 0.35$, $0.25 < z < 0.45$ and $0.35 < z < 1.0$, based on the photometric redshifts of the individual galaxies. Within each interval, we optimize the spectral template as described

---

[2] These template stellar population spectra are part of the GALAXEV package available at http://www.cida.ve/~bruzual/bc2003
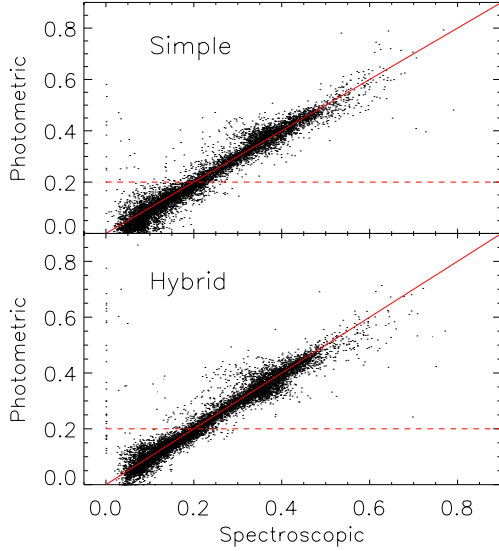[3] http://spectro.princeton.edu

Figure 4. Scatter plot showing the photometric redshift versus the spectroscopic redshift for a random 10 000 galaxies from our calibration sample. The upper panel shows the results for the simple template-fitting code of Section 3.1 and in the lower panel are the results for the hybrid code of Section 3.2. The solid (red in electronic figure) line has slope 1, while the dashed line marks the fiducial lower-redshift limit of any photometric luminous red galaxy (LRG) sample. The difficulty of estimating redshifts at $z \sim 0.4$ is evident from the increased scatter.

above. The overlapping redshift intervals provide a smooth progression in spectral type from one redshift interval to the next, as well as ensuring that the number and distribution of calibration redshifts is sufficient to constrain the colours of galaxies across a broad spectral range.

3.3 Results

We now apply the methods of the previous two sections to our calibration data set; the results are summarized in Fig. 4. Both are essentially unbiased ($|\Delta z_{\mathrm{mean}}| < 0.01$) at redshifts less that 0.5. At higher redshifts, the photometric redshifts are systematically lower than the spectroscopic redshifts by approximately 5 per cent. The scatter in both methods is approximately $\sigma \sim 0.035$, except at redshifts greater than 0.55, where the scatter grows to $\sim$0.06, caused both by increased photometric scatter and increased uncertainties in templates (caused by, for example, star formation or emission lines).

There are two noticeable features in Fig. 4 that deserve comment. The first is that the hybrid methods do significantly better than the single-template fits at low redshifts ($z < 0.15$). This is a result of the failure of the LRG selection criteria at low redshifts; a single elliptical template no longer well describes this population. This highlights an important advantage of the hybrid methods: they adjust their templates to better describe the populations.

The second feature is the increased scatter around $z \sim 0.4$, caused by an accidental degeneracy as a result of the SDSS filters. Fig. 1 shows a gap between the g and r bands at approximately 5500 Å.[4] As the 4000-Å break enters this gap at $z \sim 0.38$, the lack of

[4] This gap is partly intentional, avoiding the O I (5577 Å) night-sky emission lines. However, the filters were intended to overlap more.
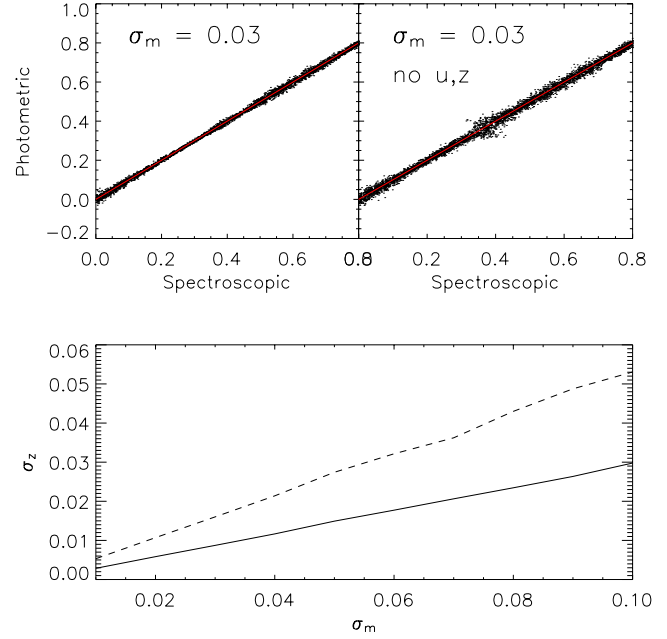

Figure 5. Simulations showing the effect of magnitude errors on the accuracy of the photometric redshifts. The upper left plot shows the reconstructed photometric redshifts for a magnitude error, $\sigma_{\mathrm{m}} = 0.03$ in all five bands, while the upper right panel has no S/N in the u and z bands and $\sigma_{\mathrm{m}} = 0.03$ in the remaining bands. The lower panel shows the redshift error induced by magnitude errors; the solid line has constant error across the bands, while the dashed line has constant error in g, r, i and zero S/N in u and z. Because the magnitude errors are independant of redshift, the redshift errors are simply computed over the entire redshift range.

coverage in either the g or r band causes a degeneracy between the strength of the 4000-Å break and its location, increasing the redshift errors.

It is useful to be able to separate the effects of template errors from photometric errors in the redshift error budget. In order to do this, we simulate galaxies by uniformly distributing them between $0 < z < 0.8$ with synthetic colours given by the template used in the method of Section 3.1. We then add errors to these synthetic fluxes, focussing on two extreme cases: uniform errors across all five bands; and zero signal-to-noise ratio (S/N) in the u and z band (i.e. infinite magnitude errors corresponding to non-detection in the u and z band, with uniform errors in the other bands). The latter case is motivated by the fact that the SDSS camera is least sensitive in the u and z bands, and because most LRGs are not detected in the u band. The results of this exercise are shown in Fig. 5. The upper panels show a realization with a (optimistic) magnitude error, $\sigma_{\mathrm{m}} \sim 0.03$. For comparison, the median S/N ($\sim 1/\sigma_{\mathrm{m}}$) for LRGs at $z \sim 0.3$ is $\sim$(2, 30, 70, 80, 30) and $\sim$(0.5, 10, 25, 36, 15) at $z \sim 0.5$. A prominent feature is the degeneracy at $z \sim 0.4$ discussed above, for the case where the u and z bands have no signal. In this case, there is no extra information that can be used to break the degeneracy between the 4000-Å break strength and its location. Also, the scatter in the photometric redshifts increases for redshifts greater than $\sim$0.35, coinciding with the 4000 Å break moving into the r band and the loss of redshift information from the $g - r$ colour. The lower panel shows how the redshift errors increase with increasing magnitude errors, again for the cases of uniform S/N in all bands, and in g, r, i with zero S/N in u and z. We also note that the redshift errors we measure are consistent with being caused principally by photometric
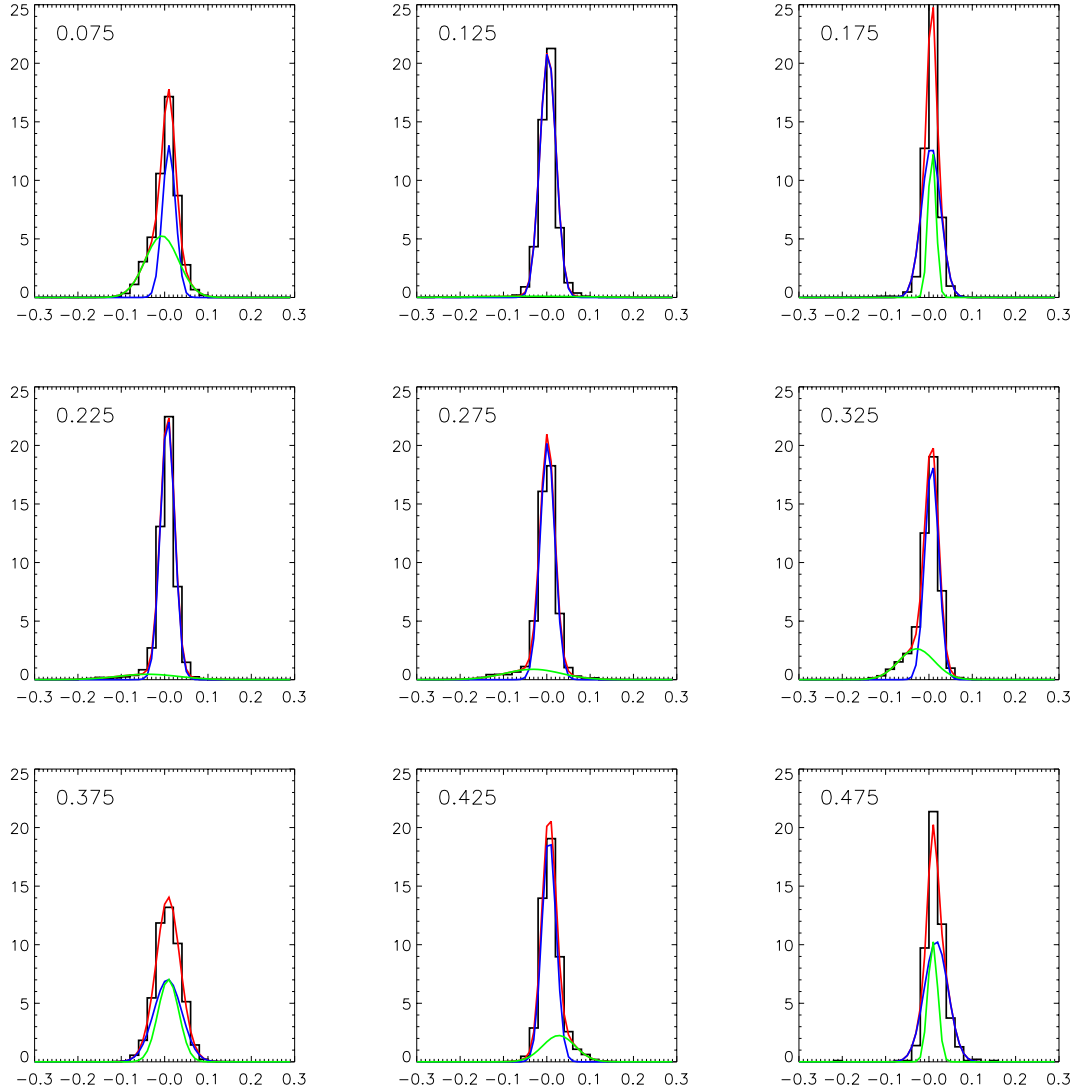
Figure 6. The double Gaussian fits to the error distribution as a function of spectroscopic redshift. The *x*-axis shows the redshift error, $z_{\mathrm{spectro}} - z_{\mathrm{photo}}$, and each panel is a redshift slice with the central redshift shown in the upper left. The histogram is the measured distribution, while the curves are the best-fitted Gaussian (both individually and summed). The data here are Sloan Digital Sky Survey (SDSS) galaxies selected using Cut I. The photometric redshifts are estimated using the method of Section 3.2.

scatter. However, the bias seen at high redshifts cannot be caused by photometric errors and suggest either errors in the template, or errors in the photometric zero-points or filter curves.

In order to parametrize the error distribution, we divide the calibration data into redshift slices of width 0.05 (except between $z = 0.6$ and 0.7, which we combine because of the small number of galaxies in that range). Within each of these ranges, we fit the error distribution with a sum of Gaussians,

$$P(\delta z = z_{\mathrm{spectro}} - z_{\mathrm{photo}}) = \frac{1}{\mathcal{N}} \sum_{i=1}^{N} b_i \exp\left[\frac{-(\delta z - m_i)^2}{2\sigma_i^2}\right], \quad (11)$$

where $\mathcal{N}$ is the normalization given by

$$\mathcal{N} = \sum_{i=1}^{N} \sqrt{2\pi} b_i \sigma_i. \quad (12)$$

We find, as shown in Figs 6 and 7, that the error distribution is well approximated by two Gaussians. The parameters of the fits for both the simple template and hybrid methods are in Tables 1 and 2, respectively. We note that the cores of these error distributions are significantly tighter than the errors mentioned above. However, the error distributions typically have long wings that are responsible for most of the measured rms errors. The discrepancies between the SDSS and SDSS–2dF samples in the overlap region are the result of a colour bias introduced by the sharp colour cuts, resulting in a bias in the redshift estimation for Cut II galaxies between $z = 0.35$ and 0.45. We therefore recommend using the SDSS distributions to $z = 0.45$ for samples constructed by combining Cut I and Cut II.

In addition to measuring the error distribution, it is useful to measure the fraction of galaxies whose redshifts are catastrophically wrong. We define a catastrophic failure as a photometric redshift that differs from the spectroscopic redshift by more than Δz_{c}, where we use $\Delta z_{\mathrm{c}} = 0.1$ and 0.2. For $\Delta z_{\mathrm{c}} = 0.1$, we have a catastrophic failure rate of 3.5 per cent for the simple template-fitting algorithm and 1.5 per cent for the hybrid algorithm. However, a large fraction of this is dominated by the underestimation of the photometric redshifts
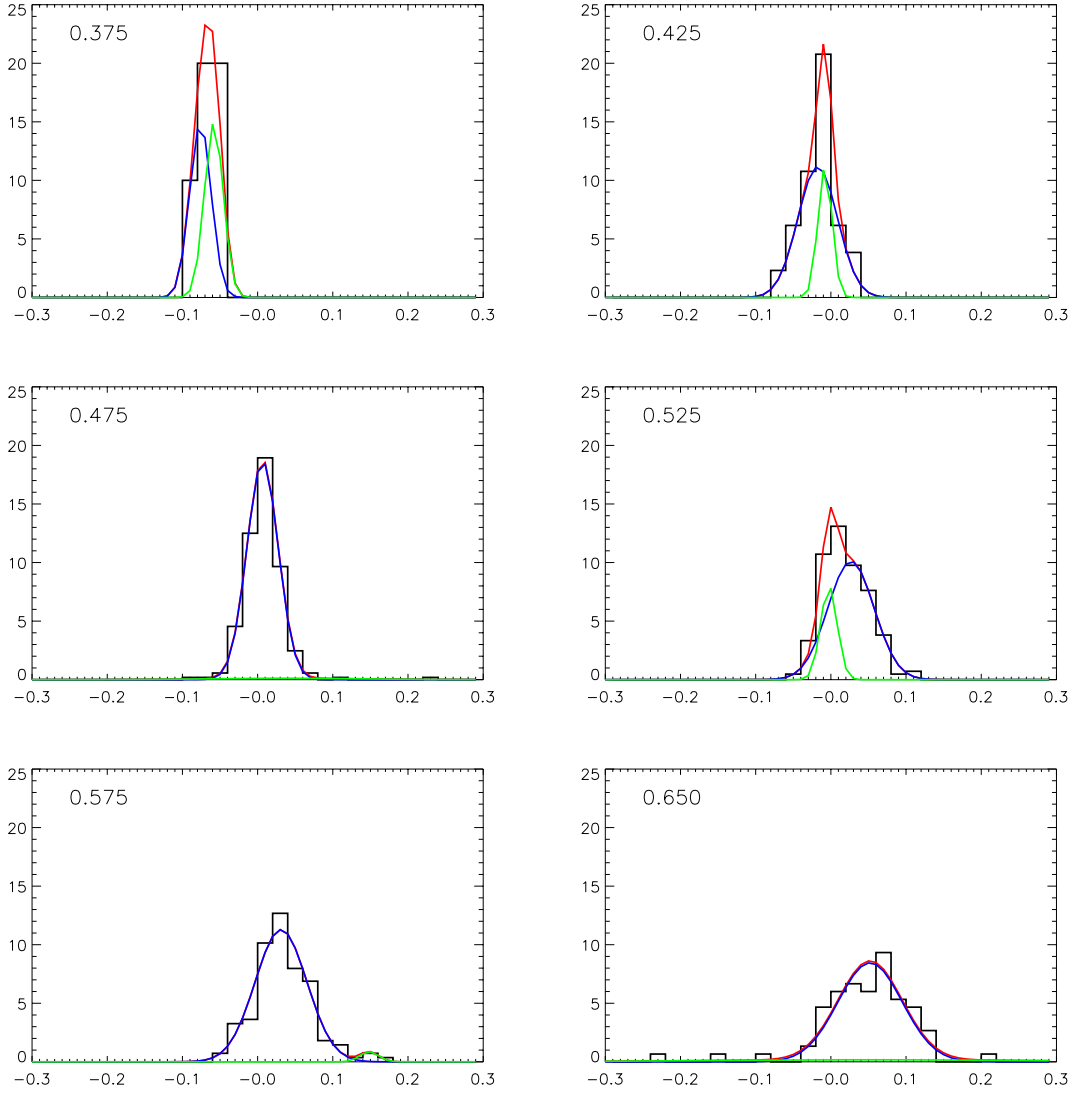
Figure 7. The same as Fig. 6 except for SDSS–2dF galaxies selected using Cut II and from redshifts 0.35 to 0.7.

at $z > 0.5$. If we increase Δz_c to 0.2, the failure rate drops to under 0.5 per cent.

4 ESTIMATING d*N*/d*z*

In the previous section, we estimated the photometric redshift error distributions as a function of the true redshift of the galaxy. With this in hand, we turn to the problem of estimating the actual redshift distribution, dN/dz of a sample of galaxies given the distribution of their photometric redshifts, $[dN/dz]_p$.

As discussed in the Section 1, the apparently trivial solution to this problem is to measure the error distribution not as a function of the true redshift, but as a function of photometric redshift. One can then add these distributions, weighted by $[dN/dz]_p$ to estimate the true redshift distribution. The problem with this approach is that the photometric redshift error distributions measured as a function of photometric redshift depend sensitively on the selection criteria of the calibration sample. If these criteria do not match those of the full sample (and in general, they will not), then dN/dz estimated using the above technique will be biased.

In order to proceed, we observe that the photometric redshift distribution is simply the convolution of the true redshift distribution with redshift errors,

$$\left[\frac{dN}{dz}\right]_p \sim \left[\frac{dN}{dz}\right] \otimes \Delta z. \tag{13}$$

If we define $\Delta(z - z_p, z)$ as the probability that a galaxy at redshift z is scattered to photometric redshift z_p, then we can write out the above more concretely:

$$\left[\frac{dN}{dz}\right]_p (z_p) = \int_0^\infty dz' \left[\frac{dN}{dz}\right] (z') \Delta\left(z' - z_p, z'\right), \tag{14}$$

where the left side has the known $[dN/dz]_p$, while the right is the unknown dN/dz. Equation (14) is a Fredholm equation of the first kind[5] and is ubiquitous throughout astronomy (Craig & Brown 1986). Unfortunately, such problems do not possess a unique solution and, moreover, are ill-conditioned. Small perturbations in the data can

---

[5] For a non-technical introduction, see Press et al. (1992), chapter 18.

Table 1. Double Gaussian fits to the photometric redshift error, $z - z_p$ as a function of z for the Sloan Digital Sky Survey (SDSS) and SDSS–2dF data. (m_1, σ_1) and (m_2, σ_2) are the mean and standard deviation of the first and second Gaussians respectively, while b is the ratio of the amplitude of the second Gaussian to the first. The photometric redshifts were computed using the method of Section 3.1. We recommend using the SDSS distributions to $z = 0.45$ and SDSS–2dF for higher redshifts.

SDSS/SDSS–2dF photometric redshift errors Single-template fitting Double Gaussian fits						
Catalogue	z_{bin}	m_1	σ_1	m_2	σ_2	b
SDSS	0.05–0.10	0.031	0.045	−0.194	0.084	0.002
SDSS	0.10–0.15	0.008	0.019	0.029	0.060	0.119
SDSS	0.15–0.20	0.010	0.015	0.013	0.052	0.059
SDSS	0.20–0.25	0.008	0.036	0.011	0.013	1.000
SDSS	0.25–0.30	0.002	0.018	−0.026	0.070	0.058
SDSS	0.30–0.35	−0.002	0.019	−0.030	0.040	0.338
SDSS	0.35–0.40	−0.004	0.028	0.000	0.000	1.000
SDSS	0.40–0.45	0.007	0.028	−0.003	0.014	1.000
SDSS	0.45–0.50	0.015	0.023	0.001	0.011	1.000
SDSS–2dF	0.35–0.40	−0.064	0.026	−0.069	0.002	1.000
SDSS–2dF	0.40–0.45	−0.019	0.022	−0.330	0.000	0.000
SDSS–2dF	0.45–0.50	0.003	0.030	0.010	0.016	1.000
SDSS–2dF	0.50–0.55	0.018	0.035	−0.278	0.000	0.000
SDSS–2dF	0.55–0.60	0.023	0.049	0.066	0.113	0.021
SDSS–2dF	0.60–0.70	0.039	0.047	0.011	0.093	0.423

Table 2. The same as Table 1 except that the photometric redshifts were computed using the methods of Section 3.2.

SDSS/SDSS–2dF photometric redshift errors Hybrid method Double Gaussian fits						
Catalogue	z_{bin}	m_1	σ_1	m_2	σ_2	b
SDSS	0.05–0.10	0.011	0.014	−0.005	0.036	0.491
SDSS	0.10–0.15	0.001	0.027	0.005	0.011	1.000
SDSS	0.15–0.20	0.005	0.022	0.007	0.009	1.000
SDSS	0.20–0.25	0.008	0.016	−0.040	0.060	0.025
SDSS	0.25–0.30	0.002	0.017	−0.028	0.063	0.050
SDSS	0.30–0.35	0.006	0.016	−0.029	0.041	0.136
SDSS	0.35–0.40	0.006	0.032	0.011	0.026	1.000
SDSS	0.40–0.45	0.014	0.027	0.004	0.012	1.000
SDSS	0.45–0.50	0.019	0.025	0.008	0.011	1.000
SDSS–2dF	0.35–0.40	−0.076	0.014	−0.058	0.012	1.000
SDSS–2dF	0.40–0.45	−0.017	0.026	−0.008	0.009	1.000
SDSS–2dF	0.45–0.50	0.007	0.021	0.037	0.098	0.006
SDSS–2dF	0.50–0.55	0.027	0.031	−0.003	0.011	0.792
SDSS–2dF	0.55–0.60	0.031	0.034	0.148	0.013	0.075
SDSS–2dF	0.60–0.70	0.052	0.043	0.053	0.227	0.020

produce solutions that are arbitrarily different. This is not surprising, given that equation (14) describes a smoothing operator that generically loses information, implying that the solution will in general require incorporating some prior knowledge about $\mathrm{d}N/\mathrm{d}z$.

4.1 Discretization and the classical solution

We begin by approximating $[\mathrm{d}N/\mathrm{d}z]_\mathrm{p}$ as a step-wise constant function measured in n bins, $[z_\mathrm{p}^i, z_\mathrm{p}^{i+1})$ with $i = 0, \ldots n - 1$, and $\mathrm{d}N/\mathrm{d}z$ in m bins, $[z^j, z^{j+1})$ where $j = 0, \ldots, m - 1$. Substituting into

equation (14), we obtain

$$\left[\frac{\mathrm{d}N}{\mathrm{d}z}\right]_{\mathrm{p},i} = \mathbf{A}_{ij}\left[\frac{\mathrm{d}N}{\mathrm{d}z}\right]_j,\tag{15}$$

where we assume the Einstein summation convention. The response matrix $\mathbf{A}_{ij}$ is given by

$$\mathbf{A}_{ij} = \frac{1}{z_\mathrm{p}^{i+1} - z_\mathrm{p}^i}\int_{z_\mathrm{p}^i}^{z_\mathrm{p}^{i+1}}\mathrm{d}z_\mathrm{p}'\int_{z^j}^{z^{j+1}}\mathrm{d}z'\Delta(z' - z_\mathrm{p}', z').\tag{16}$$

For the specific case where Δ can be described by a sum of N Gaussians (equation 11), one can perform one of the integrals explicitly to obtain

$$\mathbf{A}_{ij} = \frac{1}{z_\mathrm{p}^{i+1} - z_\mathrm{p}^i}\int_{z_\mathrm{p}^i}^{z_\mathrm{p}^{i+1}}\mathrm{d}z_\mathrm{p}'\sqrt{\frac{\pi}{2}}\sum_{k=1}^{N}b_k\sigma_k$$
$$\times\left[f(\bar{z}_k, z^{j+1}, \sigma_k, z_\mathrm{p}') - f(\bar{z}_k, z^j, \sigma_k, z_\mathrm{p}')\right],\tag{17}$$

where we define

$$f(\bar{z}, z^j, \sigma, z) = \mathrm{erf}\left(\frac{|z^j - \bar{z} - z|}{\sqrt{2}\sigma}\right)\mathrm{sgn}\left(\frac{z^j - \bar{z} - z}{\sqrt{2}\sigma}\right),\tag{18}$$

where sgn is the sign operator and erf is the error function. Note that discretizing the problem has recast an integral equation (equation 14) into a matrix problem (equation 15), albeit with a non-square matrix. We can obtain a solution to this problem by singular value decomposition (SVD; Press et al. 1992). We denote this the classical solution because we do not explicitly use any prior information about $\mathrm{d}N/\mathrm{d}z$.

In order to understand the behaviour of the classical solution, we test it on simulations of the photometric redshift distribution. We start by distributing galaxies randomly in redshift between $z = 0$ and $z = 1$ according to

$$\frac{\mathrm{d}N}{\mathrm{d}z} = \frac{z^2}{1 + \exp(20z - 14)}.\tag{19}$$

This distribution initially grows as z^2, and is exponentially cut off at $z \sim 0.6$ and approximates a volume limited distribution with a magnitude limit at high redshifts. Random redshift errors, using the model of Table 1, are added to obtain photometric redshifts, z_p. For redshifts greater than 0.7, the errors are sampled from a Gaussian whose mean and width are obtained by linearly extrapolating the errors from Table 1. Finally, we restrict to galaxies with $z_\mathrm{p} \in [0.1, 0.7]$. An example of the true and photometric redshift distributions is shown in Fig. 8.

The photometric redshift distribution is then discretized into n bins, $[\mathrm{d}N/\mathrm{d}z]_{\mathrm{p},i}$. We present results for $n = 10$ ($\Delta z = 0.06$) and $n = 15$ ($\Delta z = 0.04$). The estimated $\mathrm{d}N/\mathrm{d}z$ is likewise parametrized as a piecewise constant function from $z = 0$ to 1 with a step width of $\Delta z = 0.05$. Using these parametrizations, we construct the response matrix $\mathbf{A}_{ij}$ (equation 16) and solve for $\mathrm{d}N/\mathrm{d}z$ using equation (15). For the parameters considered here (and indeed, for generic choices), this is an underdetermined linear system. We solve it using SVD and back-substitution (Press et al. 1992), setting singular values $< 10^{-5}$ to zero. Fig. 8 shows the estimated $\mathrm{d}N/\mathrm{d}z$ averaged over 50 simulations and compares it to the true redshift distribution.

The first observation is that the classical solution reconstructs the redshift distribution accurately for certain choices of discretizations and, in particular, for discretizations of the photometric distributions with step sizes approximately the width of the photometric redshift errors. The largest errors are for $z \geqslant 0.9$ that result from the fact that $\mathrm{d}N/\mathrm{d}z$ is almost completely unconstrained at these redshifts by
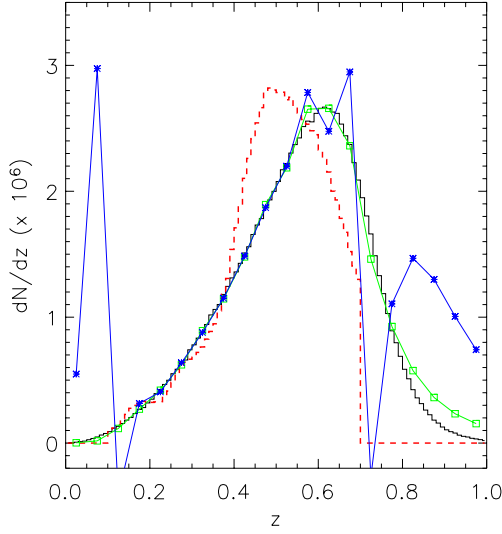
Figure 8. Results of simulations of the classical solution of the redshift inversion problem. The solid (black) histogram is the true redshift distribution, while the broken (red in electronic figure) histogram shows the photometric redshift distribution. The connected boxes (green in electronic figure) and stars (blue in electronic figure) show the reconstructed redshift distributions for different discretizations (10 and 15 bins, respectively) of the photometric redshift distributions. In both cases, the reconstructed distribution is parametrized by 20 step functions.

$[dN/dz]_p$ as only 6 per cent of objects at these redshifts scatter to $z \leqslant 0.7$.

We also observe that as we increase the resolution of $[dN/dz]_p$, the reconstruction goes unstable, ringing at the edges of the photometric redshift catalogue. Note that the reconstructions in Fig. 8 are averages and the instabilities in a single reconstruction are significantly larger.

This behaviour has a simple, intuitive explanation. The effect of photometric redshift errors is to smooth away the high-frequency $[(\text{redshift error})^{-1}]$ components in dN/dz. However, $[dN/dz]_p$ has high frequencies as a result of noise in the data and these induce large oscillations in the reconstruction. To be more quantitative, we start with a simplified model of the photometric errors,

$$\Delta(z - z_p, z) \propto \exp\left[\frac{-(z - z_p)^2}{2\sigma^2}\right]. \tag{20}$$

A component of dN/dz with frequency k will be attenuated by a factor of

$$\int dz \exp\left[\frac{-(z - z_p)^2}{2\sigma^2}\right] e^{ikz} \propto \exp\left(-\frac{k^2\sigma^2}{2}\right). \tag{21}$$

However, $[dN/dz]_p$ has a Poisson noise component that tends to a constant at high frequencies. Therefore, the inversion excites high-frequency modes in the reconstruction with amplitude $\propto \exp(k^2\sigma^2/2)$. Equation (21) also implies that this becomes significant for modes with $k > 1/\sigma$, agreeing with our intuitive picture.

The effect of the discretization step size on the the stability of the classical solution is now clear: discretization cuts off frequencies higher than $\sim 1/\delta z$ where δz is the step size, filtering out the problematic modes. This also suggests that the ideal discretizations have $\delta z \sim \sigma$, as demonstrated in our simulations.

4.2 Regularization

We would like to modify the classical solution so that it becomes less sensitive to the inversion instability discussed in the previous section. In order to do so, it is useful to rephrase the classical solution as a minimization problem.[6] If we define the energy functional,

$$E_0 = \left| \mathbf{A}_{ij} \left[\frac{dN}{dz}\right]_j - \left[\frac{dN}{dz}\right]_{p,i} \right|^2, \tag{22}$$

then the classical solution is the value of dN/dz that minimizes E_0. Given this description, it is trivial to include a penalty function that imposes smoothness on the reconstructed function,

$$E = E_0 + \lambda P, \tag{23}$$

where P is the penalty function and λ adjusts the relative weight of P in the minimization of E.[7] There are number of possible choices for the P that would impose smoothness; we use the forward difference operator,

$$P = \sum_{j=0}^{m-1} \left(\left[\frac{dN}{dz}\right]_{j+1} - \left[\frac{dN}{dz}\right]_j\right)^2. \tag{24}$$

There remains the problem of choosing an appropriate value for λ. Unfortunately, there is no a general method for choosing an optimal value. The best that we can do is to define a general merit function that objectively selects an appropriate range for λ. Based on the discussion in Craig & Brown (1986), we use

$$\Xi^2 = \frac{1}{n}\sum_{i=0}^{n-1}\left[\left(\mathbf{A}_{ij}\left[\frac{dN}{dz}\right]_{\text{av},j} - \left[\frac{dN}{dz}\right]_{p,i}\right)\right]^2 + \frac{1}{m}\sum_{j=0}^{m-1}\left\langle\left(\left[\frac{dN}{dz}\right]_j - \left[\frac{dN}{dz}\right]_{\text{av},j}\right)^2\right\rangle, \tag{25}$$

where the average reconstruction $[dN/dz]_{\text{av},j}$ is estimated either from simulations or bootstrap resampling. The first term is a measure of how well the reconstructed dN/dz reproduces the observed $[dN/dz]_p$; this term is minimized[8] as $\lambda \to 0$ and increases with increasing λ. The second term, the error in the reconstruction, measures its stability to the presence of noise in the data. As $\lambda \to \infty$, the penalty function dominates the minimization and the reconstruction is the most stable. As λ decreases, the reconstruction is more sensitive to noise in the data, increasing this term. Choosing a value of λ near[9] the minimum of Ξ^2 picks a compromise between an accurate and stable reconstruction.

In order to test this method, we return to the simulations of the previous section. Because the regularization removes the sensitivity to the discretization of the photometric distribution, we discretize $[dN/dz]_p$ into 50 bins of thickness $\Delta z = 0.012$. The estimated redshift distribution is parametrized by 40 step functions of width $\Delta z = 0.025$. Given these parameters, we must estimate the appropriate value of λ. To do this, we run 50 simulations for a given value of

---

[6] For an alternative approach to solving this problem, see Lucy (1974).

[7] This approach appears in the literature as the method of regularization, the Phillips–Twomey method, the constrained linear inversion method and Tikhonov–Miller regularization (Press et al. 1992).

[8] Assuming the generic case of an underdetermined system, $m \gg n$.

[9] We are being intentionally vague here; the precise minimum may not be the optimal choice. However, the value of Ξ^2 provides a measure of the error that one is making as we move from the minimum.
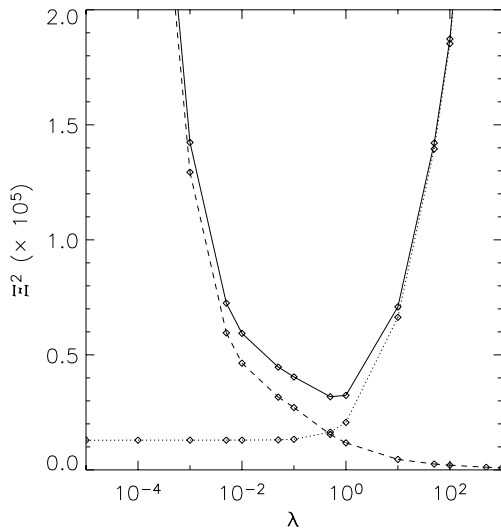
Figure 9. The value of Ξ^2 as a function of λ for the simulations discussed in the text. λ has been rescaled such that $\lambda = 1$ corresponds to equal weight being given to E_0 and P in equation (23). The dotted and dashed lines show the error and stability components of equation (25), respectively. As expected, the error term increases with increasing λ, while the stability term decreases with increasing λ. The minimum of Ξ^2 occurs near $\lambda = 0.5$.

λ to evaluate Ξ^2 and repeat this for a grid of values of λ. The results are shown in Fig. 9. We note that Ξ^2 has a well-defined minimum, with the error and stability terms demonstrating the λ dependence that we anticipated. Note that the error term does not go to zero as $\lambda \to 0$, but appears to asymptote to a non-zero constant. This is readily understood in terms of the discussion in the previous section: the measured $[dN/dz]_p$ has a high-frequency noise component that cannot be reproduced by the convolution of dN/dz with the redshift errors. It is this noise component that is responsible for the non-zero value of the error term in Ξ^2 as $\lambda \to 0$.

The upper panel of Fig. 10 shows the average of 50 reconstructions of dN/dz for values of λ near the minimum of Ξ^2. We observe that for all the values of λ considered in the figure, the reconstructions closely match the input redshift distribution for all redshifts <0.7. As before, the largest discrepancies are at high redshift because of the lack of constraint as a result of the upper photometric redshift limit of 0.7. It is also instructive to consider extreme values of λ; these are shown in the lower panel of Fig. 10. For small values of λ, the reconstructions are extremely noisy, while for large values of λ, the penalty function dominates the reconstruction. Note that the forward difference operator (equation 24) represents a constant prior, which is what we see the reconstructions approaching as $\lambda \to \infty$.

We make one cautionary observation. Based on Fig. 10, one might conclude that the best strategy for choosing λ is to preferentially choose a smaller value than that suggested by the minimum of Ξ^2. We however discourage this because, as indicated in Fig. 9, such reconstructions are very noisy. This lack of stability would result in small errors in the redshift error distribution being amplified in the reconstructions.

How many galaxies are required for the inversion? The simulations discussed above used 100 000 galaxies, similar to the expected number of photometric LRGs over the same area of sky. We have however tested the inversion on as few as 1000 galaxies and found that, for appropriate regularizations, the algorithm reconstructs the input redshift distribution. However, for small samples, the Poisson
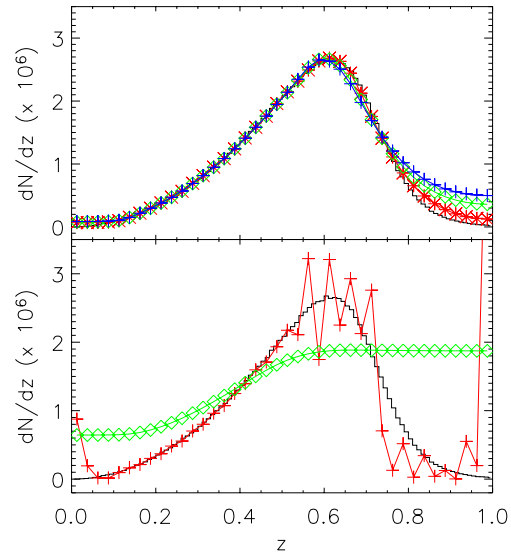

Figure 10. Regularized estimates of dN/dz for different values of the regularization parameter, λ. In both panels, the black histogram shows the true redshift distribution. The upper panel shows the reconstructions for values of λ about the minimum of Ξ^2; the stars (red in electronic figure), diamonds (green in electronic figure) and crosses (blue in electronic figure) correspond to $\lambda = 0.1, 0.5$ and 1.0, respectively. The lower panel shows the reconstructions for extreme values of λ, the crosses (red in electronic figure) and diamonds (green in electronic figure) correspond to $\lambda = 10^{-10}$ and 1000, respectively. The values of λ have been rescaled as in Fig. 9.

noise in the input photometric redshift distribution can be significant, resulting in a noisier reconstruction (for the same redshift resolution). This may be improved by smoothing the resulting reconstruction or, equivalently, reconstructing the redshift distribution on a coarser redshift grid.

There is an important generalization of this method that should be mentioned. We introduced the concept of regularization and the penalty function to cure an instability in the deconvolution as we attempted a finer resolution of the redshift distribution. Phrased differently, the deconvolution became unstable when the input had low S/N and the prior (in the form of the penalty function) compensated for this loss of information. In the cases considered in this paper, we have used a relatively weak prior; however, if one has reliable prior information (for example, the rough shape of the redshift distribution), one can easily include that information. A strong prior will allow one to obtain a solution even in the low S/N regime. We do however remind the overzealous reader that the usual caveat about strong priors does apply in this case; the method cannot recognize an incorrect prior and will get the wrong answer if such a prior is heavily weighted.

4.3 Application to SDSS data

Before applying this algorithm to a photometric sample, we test it against the Cut I calibration data set described in Section 2.3. The results are shown in Fig. 11. The reconstructed redshift distribution correctly captures all the broad features of the true redshift distribution, including correcting for the bias at $z \sim 0.1$ and sharpening the dip at $z \sim 0.25$. Fig. 11 also highlights the inability of this method to reconstruct sharp features because these are disfavoured by the smoothness prior we impose; the inversion works best for broad features. It is worth emphasizing that most sharp features (including the feature at $z \sim 0.075$) are spurious (e.g. binning artefacts). However,
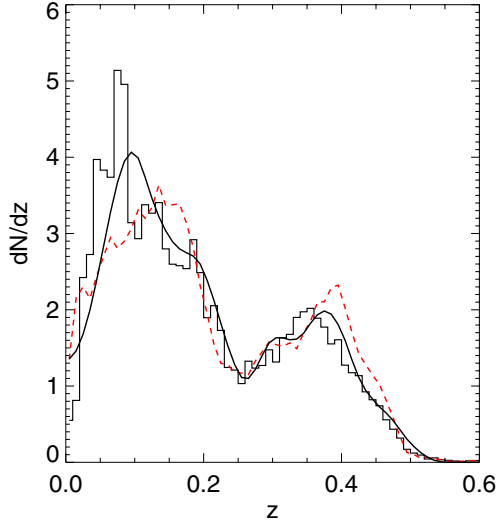
Figure 11. Regularized estimate (solid, black) of dN/dz for the Cut I calibration data (histogram). The input photometric redshift distribution is the dashed (red in electronic figure) line. dN/dz is normalized to integrate to unity.

if a sharp feature is physically expected in the distribution, the prior must be adjusted to allow for this.

We conclude this discussion by applying the above algorithm to the SDSS photometric data. A detailed discussion of the construction and properties of the SDSS photometric LRG sample will be presented elsewhere; briefly, the sample is constructed by applying the photometric selection criteria (Cut I and Cut II, see Section 2.3) to objects classified as galaxies by the photometric pipeline. We then estimate a photometric redshift for each of the selected objects using the simple template-fitting code of Section 3.1; however, the results are insensitive to the choice of algorithm. The photometric redshift distribution is shown in Fig. 12.
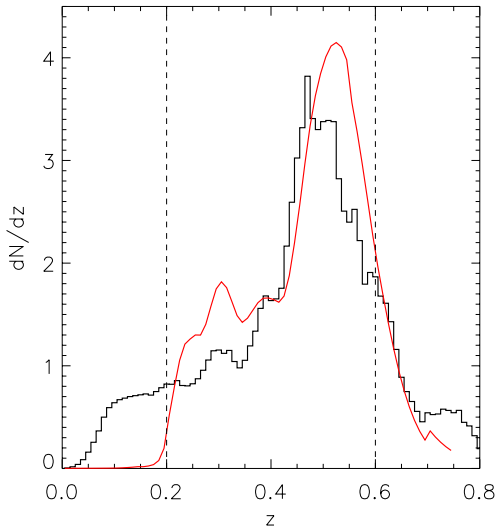
Figure 12. Regularized estimate (solid, red in electronic figure) of dN/dz for the luminous red galaxy (LRG) sample culled from the Sloan Digital Sky Survey (SDSS) photometric data, compared with the photometric redshift distribution (histogram, black). The redshift distribution is for galaxies with $0.2 < z_{photo} < 0.6$, indicated by the vertical dashed lines. As before, dN/dz is normalized to integrate to unity.

One feature of this distribution that deserves some explanation is the bump in the number of galaxies at $z \sim 0.7$. This is inconsistent with being the same population of LRGs selected with an apparent magnitude cut. It is unlikely that these are a different population at $z \sim 0.7$, as they would have to be a significantly brighter population than the LRGs, which only appeared at high redshifts. A more likely explanation is that these are faint galaxies at lower redshifts scattered to high redshifts by photometric errors. This is more likely, in light of the fact that these galaxies have $i \sim 20$, $g - r \sim 2$ and $r - i \sim 1$, giving them $r \sim 21$ and $g \sim 23$. This is at the very edge (or beyond) the photometric completeness of SDSS and the measurements will have significant photometric errors ($\sim$ tenths of a magnitude). Given such photometric errors, it is likely for the more numerous low-redshift galaxies to be scattered into the LRG colour space. Furthermore, the spectral templates that we use are not well constrained by observations for redshifts >0.7. To avoid the complications of correcting for such contamination, we restrict our catalogue to $z_{photo} < 0.6$. Similarly, as discussed earlier, the photometric selection breaks down at low redshifts and so we impose a lower-redshift cut-off of $z_{photo} > 0.2$. Note that this lower-redshift cut is imposed only to select a uniform sample; the inversion must be (and is) performed at all redshifts. However, the small photometric redshift error at these redshifts minimizes the contamination from these galaxies, effectively truncating the inverted distribution at $z \sim 0.2$.

We can now apply the inversion algorithm to estimate the true redshift distribution, using the error distributions measured in Section 3.3. The merit function, equation (25), is computed by bootstrap resampling the actual catalogue; the measured Ξ^2 has a form similar to Fig. 9. Using the value of the regularization parameter, λ, obtained from Ξ^2, we show the estimated redshift distribution for galaxies with $0.2 < z_{photo} < 0.6$ in Fig. 12. The underestimation of the photometric redshifts at high redshifts is immediately apparent from the comparison of the two distributions. The bumps from $z = 0.3$ to 0.4 are a residual artefact of the inversion. These vanish for higher values of λ and become stronger for lower regularizations, but are more unstable. The value of λ used is a balance between this stability and accuracy, as intended.

5 DISCUSSION

As we discussed in the Section 1, constructing a photometric redshift catalogue involves three steps: photometrically selecting a sample for which accurate photometric redshifts can be obtained; measuring the photometric redshift error distribution for the resulting sample; and estimating the true redshift distribution. This paper describes all stages of this process.

(i) We describe the selection of a sample of LRGs using the SDSS photometric sample. These galaxies are typically old elliptical systems with strong 4000-Å breaks in their continua. The shifting of this feature through the SDSS filters makes accurate photometric redshifts possible.

(ii) We measure the error distribution of this sample by comparing photometric and spectroscopic redshifts for a calibration subsample of galaxies culled from the SDSS and SDSS–2dF spectroscopic catalogues. The scatter in the redshifts is approximately $\sigma \sim 0.03$ at redshifts less than 0.55, and increases at higher redshifts as a result of increased photometric errors and uncertainties in the templates.

(iii) The accuracy of the photometric redshifts is similar for the two algorithms we consider: a simple template fit and a hybrid algorithm that adjusts the template to better fit the observed colour distribution.

We have specifically used the SDSS photometric sample throughout this paper, both as an example and for its intrinsic interest. However, we emphasize that the entire process that we describe can be reconstructed for any multicolour imaging survey with appropriate filters.

Using such a photometric redshift sample requires knowing the conditional probability that a galaxy with a photometric redshift, z_{photo}, has a true redshift, $z_{spectro}$. Given the redshift error distribution, this conditional probability can be readily estimated using the Bayes theorem if the true underlying redshift distribution is known. Using the fact that the photometric redshift distribution is the true redshift distribution convolved with redshift errors, we presented a method to deconvolve the errors to estimate the redshift distribution. This method is ill-conditioned and, therefore, we use a prior on the smoothness of the redshift distribution to regularize the deconvolution. We have calibrated the relative weight of this prior by measuring the accuracy and stability of the recovered redshift distributions, and we proposed a general merit function that objectively determines this weight.

We conclude with a few comments about this algorithm.

(i) The particulars of the sample selection are encoded into the photometric error distribution. The method is therefore completely general, and applicable to any combination of colour selections and photometric redshift cuts.

(ii) The accuracy of the recovered redshift distribution is determined by the accuracy of the input error distributions. Therefore, it is essential that the calibration data used to measure the error distribution correspond as closely as possible to the actual data, both in photometric depth and accuracy. One can attempt to extrapolate these distributions to fainter magnitudes or measure them from simulations, but with the caveat that the actual distributions may be very different from these and that the reconstruction could potentially be sensitive to these errors. We emphasize that this limitation is not unique to this method, but affects all analyses that use photometric redshifts.

(iii) The deconvolution algorithm is formally applicable to arbitrary error distributions. However, for complex error distributions (e.g. multiply peaked distributions), multiple solutions may exist and there is no guarantee that the method will converge to the correct solution. This problem is avoided here by the use of photometric pre-selection; in general, it could also be prevented by the use of priors in the photometric redshift estimation. We strongly recommend using one of these methods to break photometric redshift degeneracies.

(iv) An advantage to this method is that the calibration data need not sample the redshift range of interest in the same manner as the photometric data. It suffices that it samples the entire range well enough to measure the error distributions. This allows the use of calibration sets from heterogeneous samples, as was done in this paper.

(v) The inversion algorithm presented in this paper presents an alternative to iterative deconvolution algorithms (Lucy 1974; Brodwin et al. 2003). As emphasized by Lucy (1974), the two methods have very different mathematical philosophies; iterative methods treat the problem as one in statistical estimation, while the philosophy in this paper derives from the theory of integral equations. However, in the high S/N regime, both methods will produce similar results and there is little to distinguish between the two. For low S/N, the deconvolution problem may not possess a solution and iterative methods may not converge. In these cases, the algorithm presented in this paper transparently allows the inclusion of exter-

nal information as part of the penalty function to yield a meaningful solution. In cases where one possesses reliable prior information, one can then refine that information to yield a better solution.

REFERENCES

Abazajian K. et al., 2003, AJ, 126, 2081
Abazajian K. et al., 2004, AJ, 128, 502
Benítez N., 2000, ApJ, 536, 571
Blanton M. R., Lin H., Lupton R. H., Maley F. M., Young N., Zehavi I., Loveday J., 2003, AJ, 125, 2276
Bolzonella M., Miralles J.-M., Pelló R., 2000, A&A, 363, 476
Brodwin M., Lilly S. J., Porciani C., McCracken H. J., Fevre O. L., Foucaud S., Crampton D., Mellier Y., 2003, astro-ph/0310038
Brunner R. J., Connolly A. J., Szalay A. S., 1999, ApJ, 516, 563
Bruzual G., Charlot S., 2003, MNRAS, 344, 1000
Budavári T., Szalay A. S., Connolly A. J., Csabai I., Dickinson M., 2000, AJ, 120, 1588
Budavári T., Szalay A. S., Csabai I., Connolly A. J., Tsvetanov Z., 2001, AJ, 121, 3266
Budavári T. et al., 2003, ApJ, 595, 59
Cannon R., Croom S., Pimbblet K. (for the SDSS–2dF Team), 2003, AAO Newsletter, 103, 8
Coleman G. D., Wu C.-C., Weedman D. W., 1980, ApJS, 43, 393
Collister A. A., Lahav O., 2004, PASP, 116, 345
Connolly A. J., Csabai I., Szalay A. S., Koo D. C., Kron R. G., Munn J. A., 1995, AJ, 110, 2655
Craig I. J. D., Brown J. C., 1986, Inverse problems in astronomy: A guide to inversion strategies for remotely sensed data. Adam Hilger Ltd, Bristol, p. 159
Csabai I., Connolly A. J., Szalay A. S., Budavári T., 2000, AJ, 119, 69
Csabai I. et al., 2003, AJ, 125, 580
Eisenstein D. J. et al., 2001, AJ, 122, 2267
Eisenstein D. J. et al., 2003, ApJ, 585, 694
Firth A. E., Lahav O., Somerville R. S., 2003, MNRAS, 339, 1195

Fukugita M., Ichikawa T., Gunn J. E., Doi M., Shimasaku K., Schneider D. P., 1996, AJ, 111, 1748

Gladders M. D., Yee H. K. C., 2000, AJ, 120, 2148

Gunn J. E. et al., 1998, AJ, 116, 3040

Gwyn S. D. J., Hartwick F. D. A., 1996, ApJ, 468, L77

Hamilton D., 1985, ApJ, 297, 371

Hogg D. W. et al., 1998, AJ, 115, 1418

Hogg D. W., Finkbeiner D. P., Schlegel D. J., Gunn J. E., 2001, AJ, 122, 2129

Koo D. C., 1985, AJ, 90, 418

Le Borgne D., Rocca-Volmerange B., 2002, A&A, 386, 446

Lewis I. J. et al., 2002, MNRAS, 333, 279

Lucy L. B., 1974, AJ, 79, 745

Nichol R. C., 2003, in Mulchaey J. S., Dressler A., Oemler A., eds, Carnegie Observatories Astrophysics Series, Vol. 3, Clusters of Galaxies: Probes of Cosmological Structure and Galaxy Evolution. Cambridge Univ. Press, Cambridge, p. 24

Oke J. B., Gunn J. E., 1983, ApJ, 266, 713

Pier J. R., Munn J. A., Hindsley R. B., Hennessy G. S., Kent S. M., Lupton R. H., Ivezić Ž., 2003, AJ, 125, 1559

Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P., 1992, Numerical recipes in FORTRAN. The art of scientific computing, 2nd edn. Cambridge Univ. Press, Cambridge

Sawicki M. J., Lin H., Yee H. K. C., 1997, AJ, 113, 1

Schneider D. P., Gunn J. E., Hoessel J. G., 1983, ApJ, 264, 337

Smith J. A. et al., 2002, AJ, 123, 2121

Stoughton C. et al., 2002, AJ, 123, 485

Strauss M. A. et al., 2002, AJ, 124, 1810

Wang Y., Bahcall N., Turner E. L., 1998, AJ, 116, 2081

Willis J. P., Hewett P. C., Warren S. J., 2001, MNRAS, 325, 1002

York D. G. et al., 2000, AJ, 120, 1579

This paper has been typeset from a TEX/LATEX file prepared by the author.